Rev 4.0 2010

Parallels™
Optimized Computing™

Parallels Optimized Computing

Parallels Account Team

- **Adrienne Boydston, Enterprise Account Manager**
  aboydston@parallels.com
  512-406-1970

- **Blake Tyra, Senior Systems Engineer**
  blake@parallels.com
  703-310-9928

- **Michael Stump, Senior Account Executive**
  mstump@parallels.com
  703-995-6205

Agenda

- **Who is Parallels?**
- **Overview of Virtualization Technologies**
- **Parallels Containers Overview**
- **Linux Workload Performance Benchmarks**
- **Summary of Parallels Elastic IT Platform**
- **Parallels Virtual Desktop Infrastructure**
- **Q & A**

Who is Parallels?

- **Founded in 1999**
- **Over 650 Employees worldwide**
- **500+ Partners, including Microsoft, Apple, Intel, AMD, Dell, HP & IBM**

  - Currently, we maintain 85% market share in the Service Provider industry

  - Parallels currently ranks 9th of all organizations globally making contributions to the Linux Kernel

Parallels Elastic IT for Linux

Parallels is bringing the same platform we developed for the Hosting industry to the Linux Community

- Provides full-featured management tools that are **infinitely expandable** through a documented Open API (XML-RPC).
- Parallels elastic IT solution enables the creation of an on-demand IT infrastructure within the datacenter to provide **on-the-fly resizable compute capacity**
- Virtualize **high performance** Linux workloads including Java, databases and computational tasks that demand near native I/O and network speeds
- **Consolidate** Linux servers to maximize the utilization of hardware and minimize system administration

Virtualization Technology Comparison

Parallels Containers

Native Performance Architecture

Native Optimizations

- Maintains all environment optimizations
- Maintains native time clock

Minimal Overhead

- Faster context switching than hypervisors
- Better memory utilization

Suitable for performance spikes

- Scales to the full resources of the server
- Easy to migrate instances to other servers

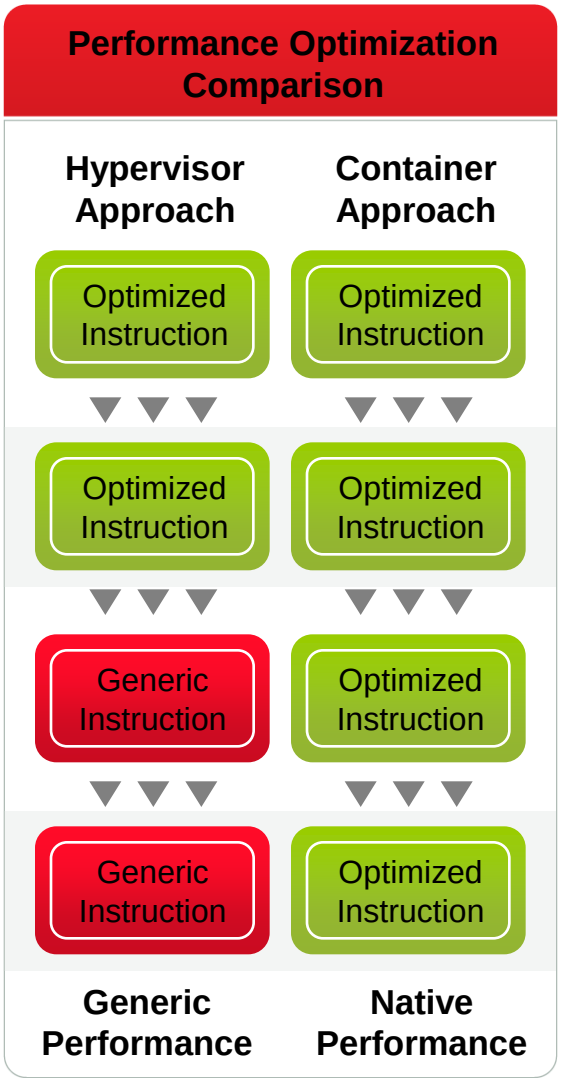

Parallels Containers: Key Benefits

High Density Consolidation

Maximize server utilization
with **negligible overhead**

Dynamic Resource Control

Change CPU, Memory, and Disk
allocations **without container reboot**

Live Migration

Move containers between physical
servers with **no service interruptions**

Mass Management

Control your physical and virtual
Infrastructure from **a single interface**

OS & Application “Templates”

Deploy Configured OS and Application Stacks En-Masse

- Provision full applications or updates to all or some of your containers

Mix-and-Match Linux Distributions on a Single Server

- Including multiple versions and editions of:
  - CentOS, Debian, Fedora Core, RedHat, SuSe, Ubuntu

Low Overhead, Minimal Initial Footprint

- New containers consume
  - 40MB disk space
  - 10MB memory
- Cache common files for substantial storage cost savings
- Share memory across containers for substantial RAM savings

Managing the Parallels Elastic IT Platform

Create and Destroy Virtual Servers in Seconds

- Quickly create / deploy containers with a few clicks
- Virtual environments are typically up and running within seconds from clicking “Create”

Control Physical & Virtual Servers from One Interface

- Web-based Parallels Virtual Automation, a multi-tenant management portal
- Enables custom integration with existing systems

Quickly Account for All Datacenter Resources

- Track resource usage on a per container basis
- Configure complex resource limitations and traffic shaping rules to match your infrastructure’s needs

Basic Functions

Provision
De-provision
Power On
Power Off
Restart
Backup
Clone
Migrate
Application
-Deployment
-Removal
-Configuration
P2V
Disk / Mem Limits
CPU Resources
File Mgmt
Network Config

Management Tools

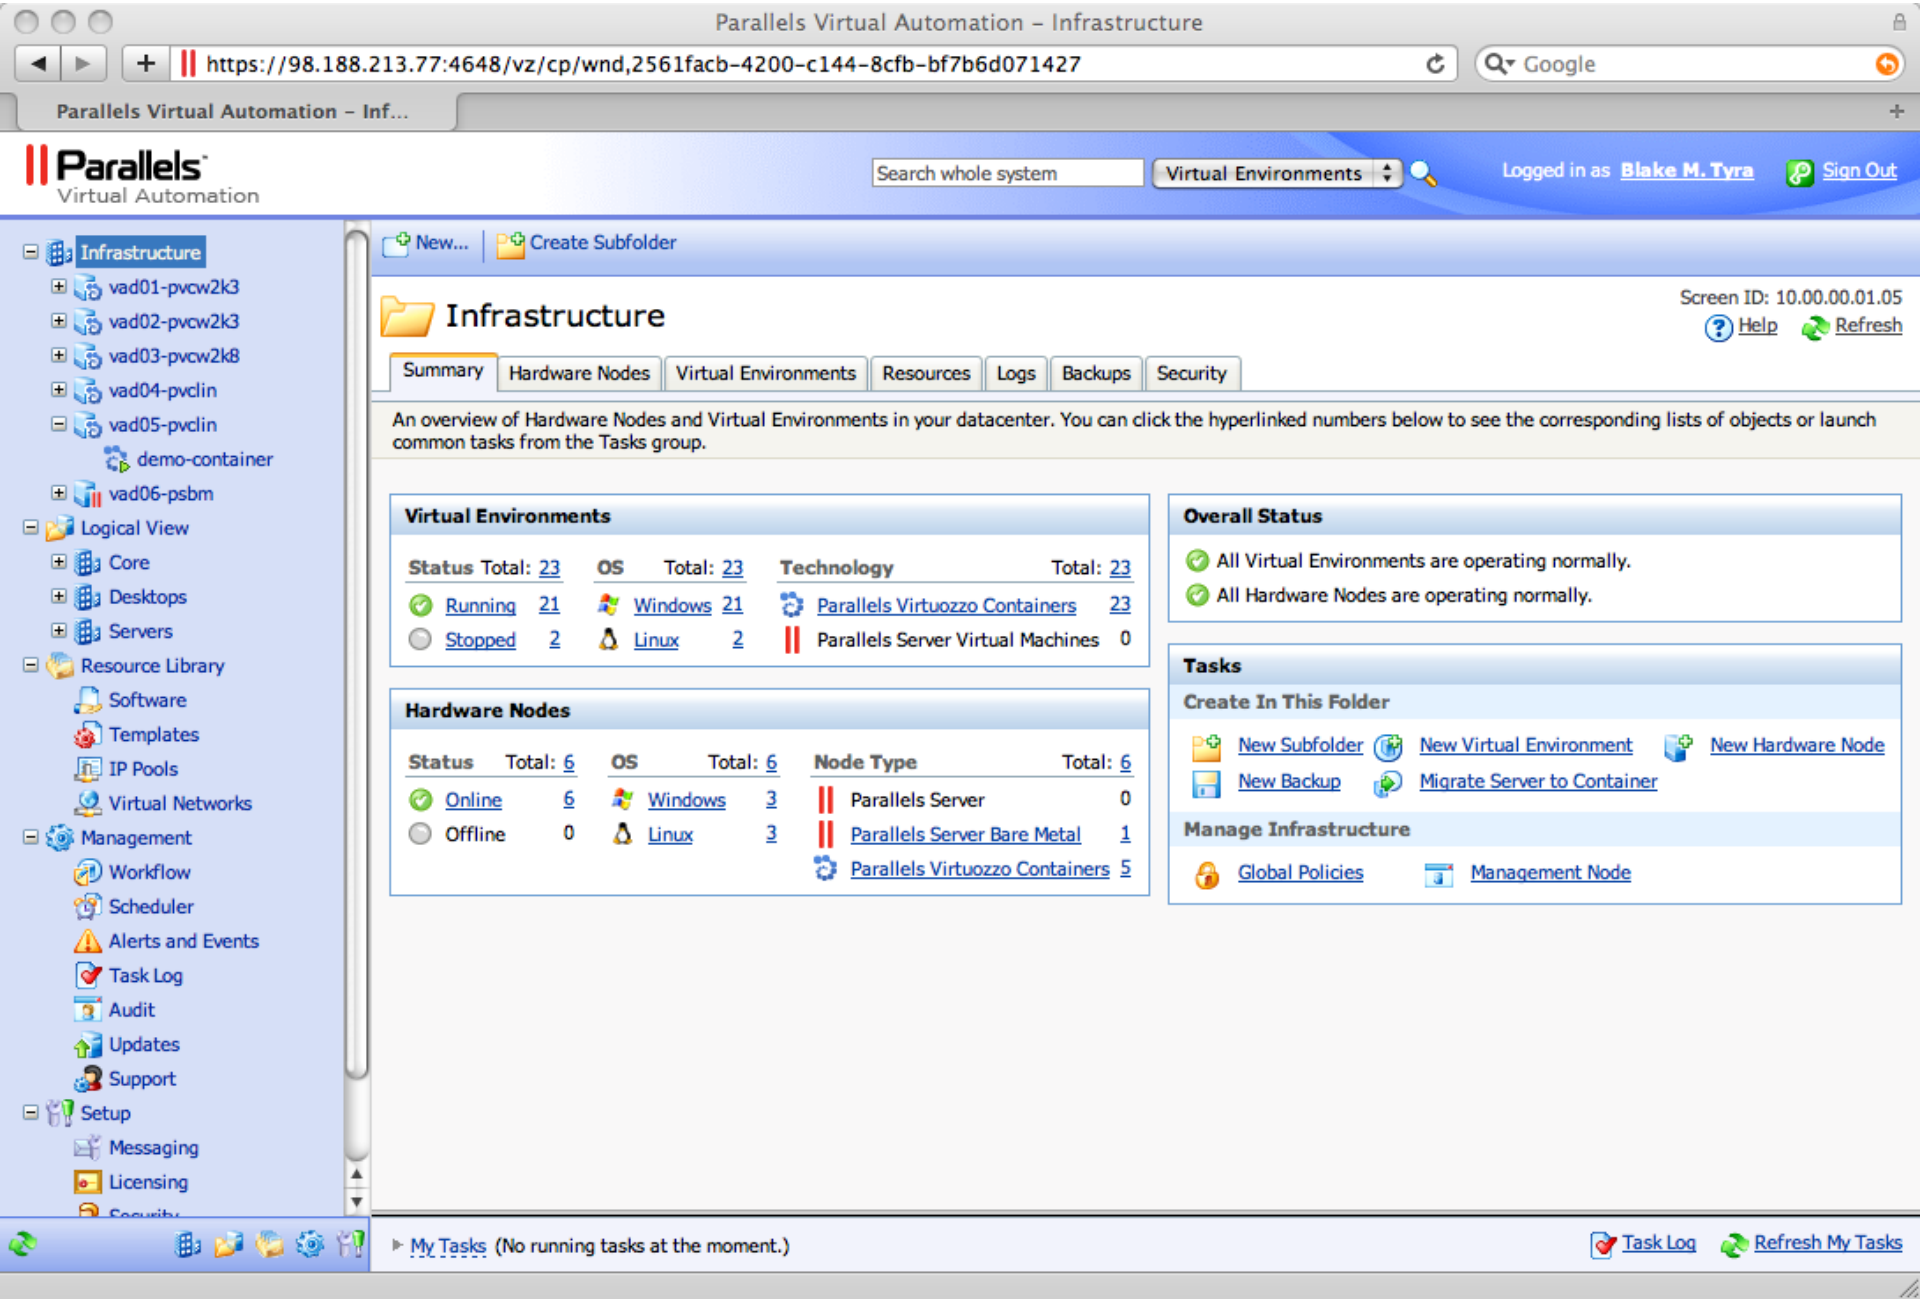

Virtual Server Snapshots

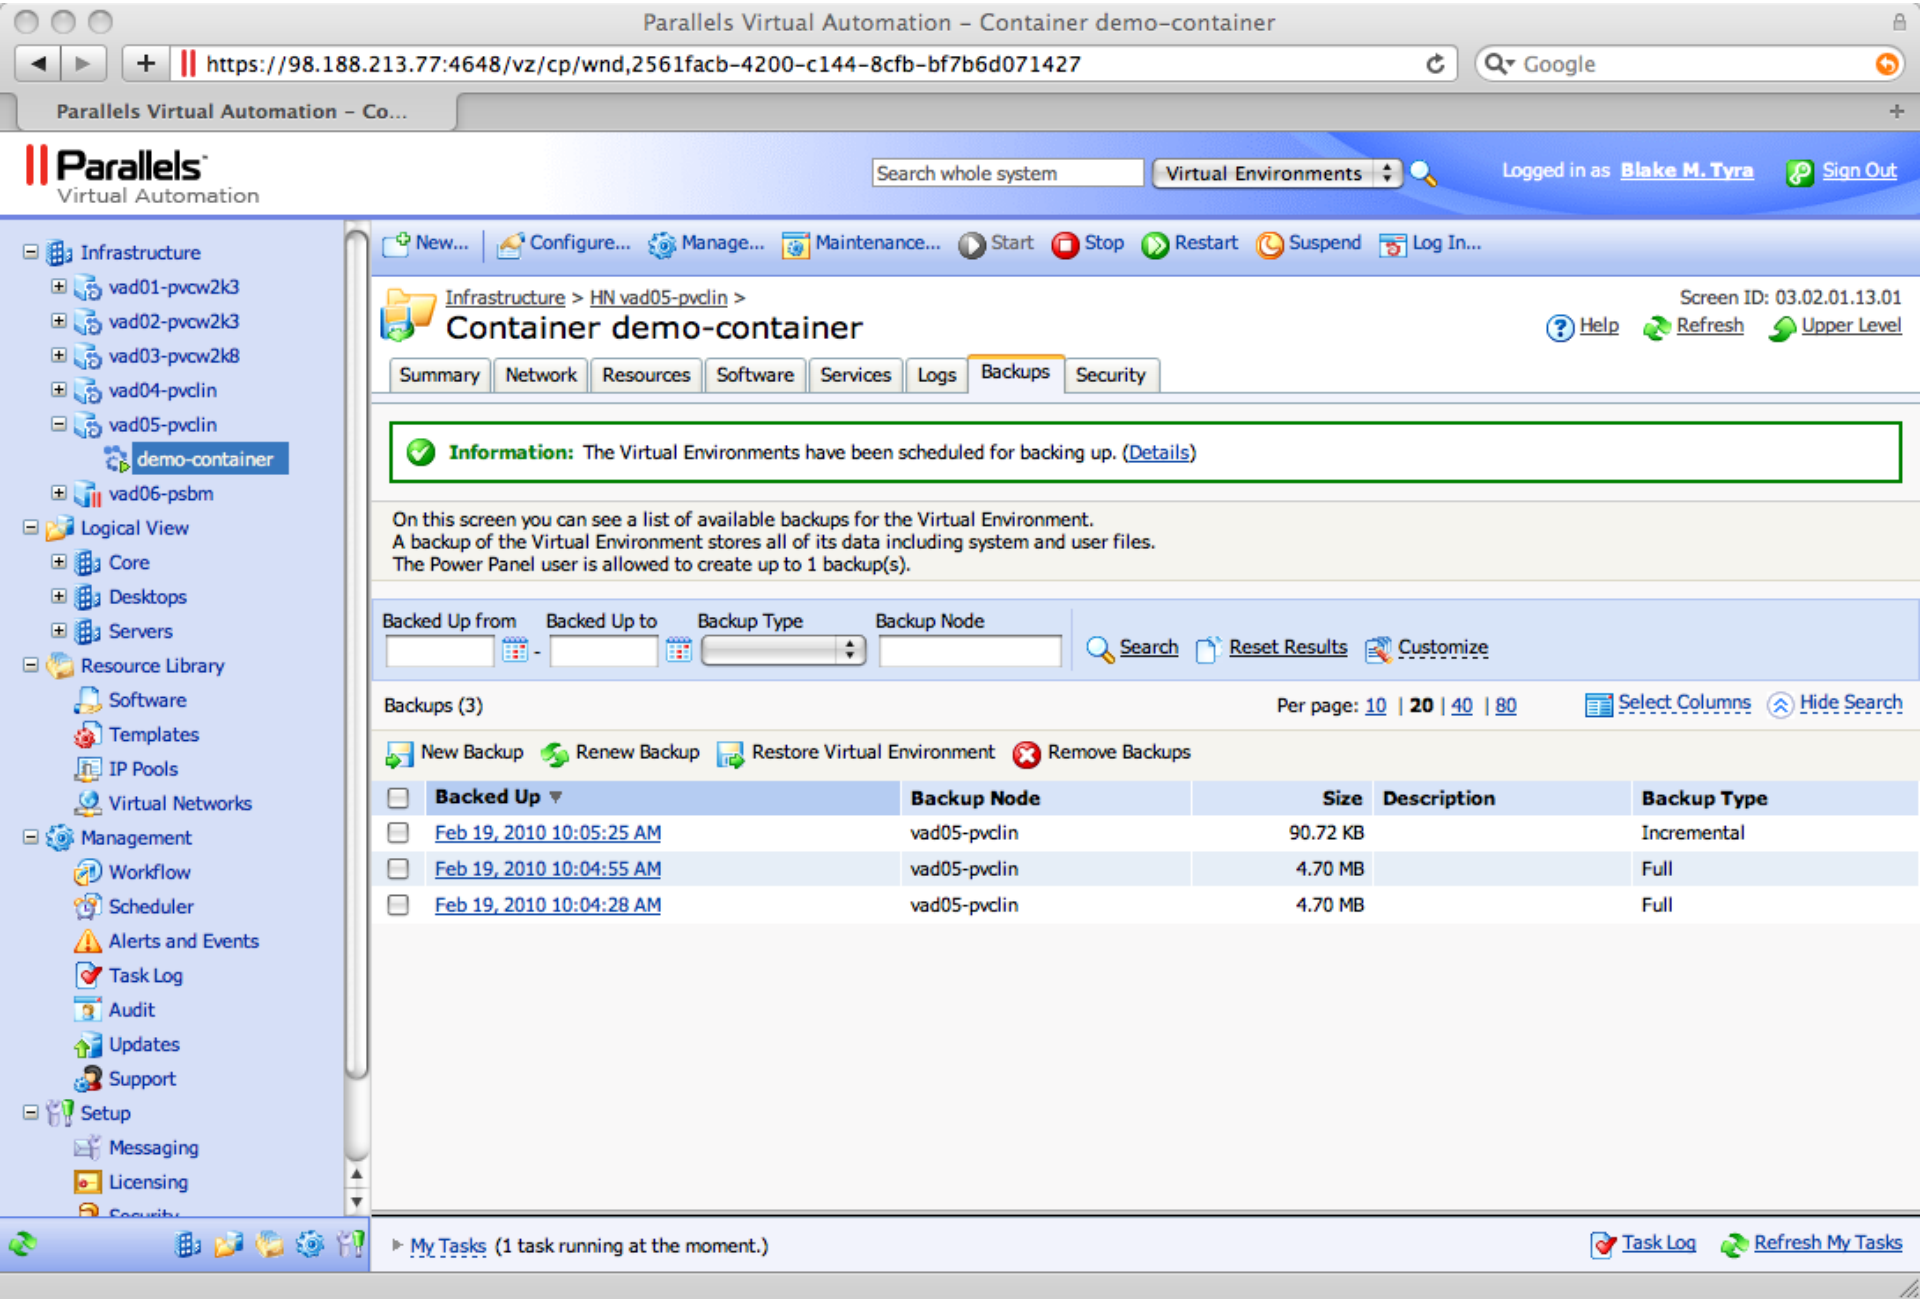

Minimal Resource Usage

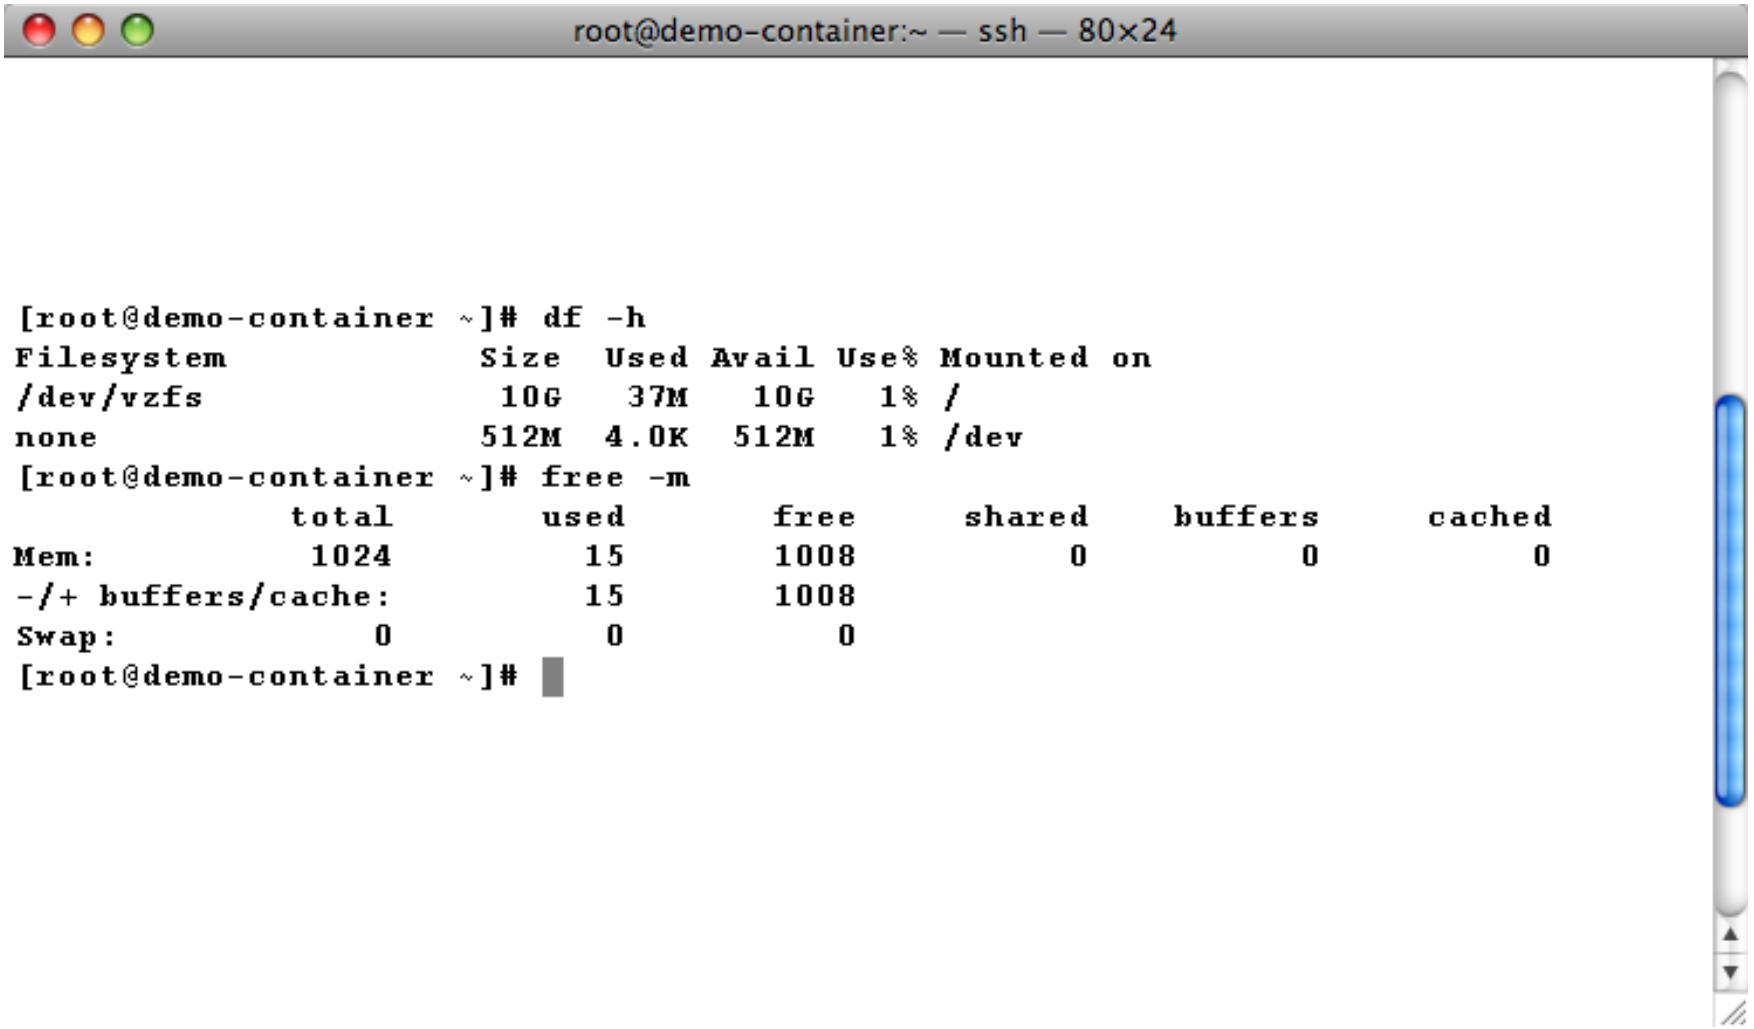

Granular Resource Controls

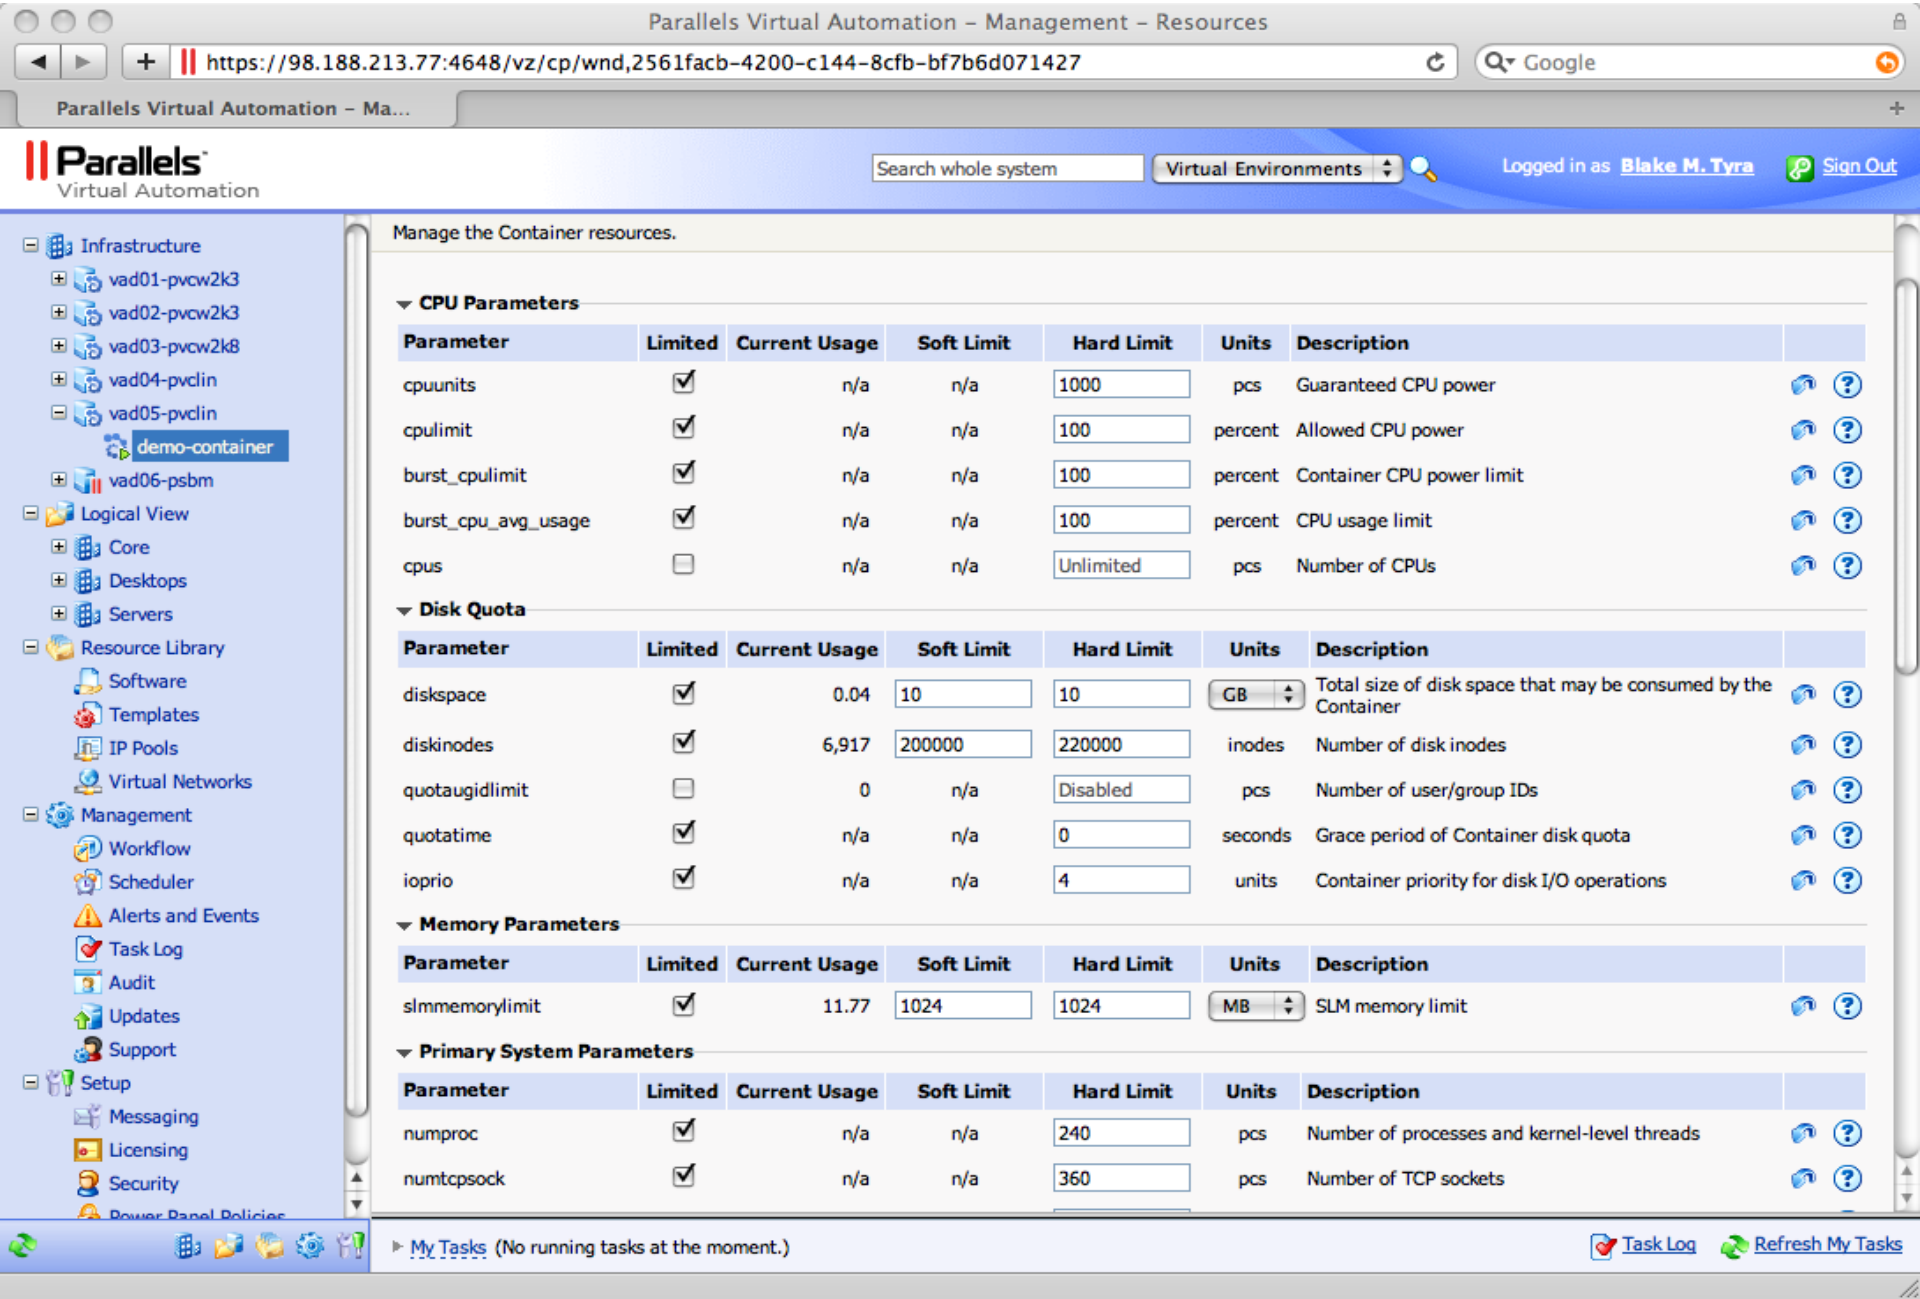

Easy Application Management

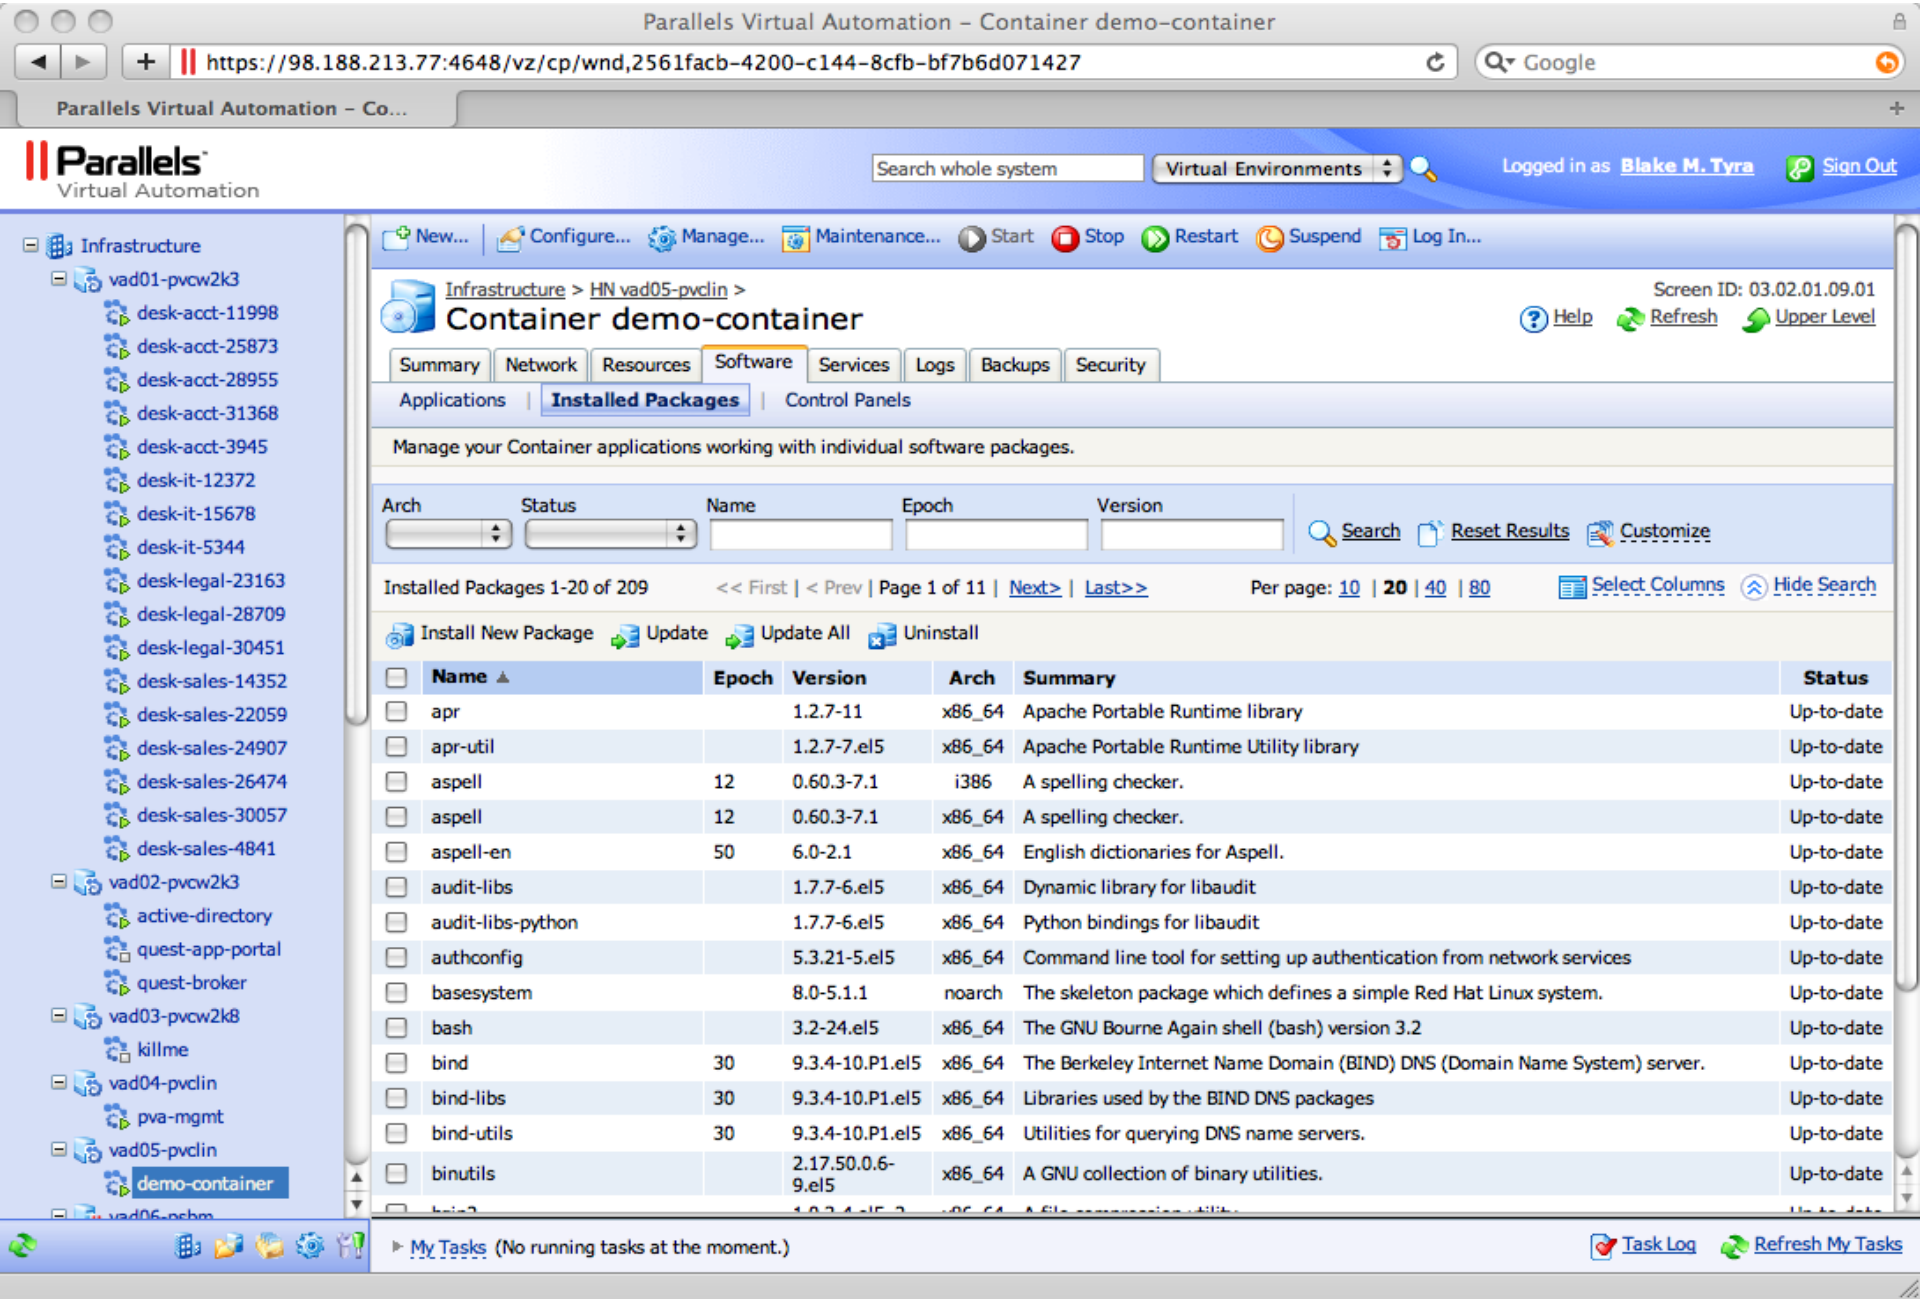

Linux Workloads: Challenge

- **Many Linux applications are single threaded (do not scale up)**
- **Deployed on many 2P X86/X64 servers (forced scale out)**
  - **Do not leverage multi-core CPU capabilities**
- **Cannot co-exist with other applications (require isolation)**
- **Results in under-utilized compute resources (typically running at ~20% load per 2P server)**

Solution: Parallels Containers for Linux

- **Customize functionality and implementation with existing systems by integrating with a fully documented XML-RPC API**
  - As flexible as an infrastructure foundation should be

- **Run a minimum of 3x more virtual environments per server as compared to hypervisors.**
  - ✓ Save on hardware, datacenter space, power consumption & administration requirements.

- **Maintain time clock consistency by using the host clock preventing the "time fade" experienced by hypervisors.**
  - ✓ Rely on the accuracy of the timestamps within applications.

- **Directly access hardware for I/O including networking & disk access to improve response times far better than going through the virtual machine manager of a hypervisor.**
  - ✓ Eliminate bottlenecks and virtualize more applications than you thought possible.

Linux Workload Performance Benchmarks

- **Measurement process**
  - simple bash script for all running CTs/VMs
    - time ssh root@ct101 “cd /root/kernel; make –j 4”&
  - Copy Linux 2.6.18 kernel source to each Parallels Container/Xen VM
  - Perform kernel compilation in for varying numbers of CTs/VMs:
    - make clean; yes “” make config; make –j 4
    - -j reflects number of CPUs available
  - Collect total compilation time for all compiles and determine average
- **Average “kernels per minute”**
  - This metric is more suitable than kernel compilation time as you can compare performance of, for example, 4 and 8 VMs in terms of “how many kernels they can compile within a specific timeframe”.

Hardware / Software

Hardware

CPU	Memory	Storage
Dual CPU / Quad Core Intel Xeon 3.0 GHz	32 GB RAM (16x2), 667 MHz	4xSAS 15K RPM Symbios Logic MegaRAID SAS 1078

Software

Host OS	Guest OS	Containers/VM config
RHEL 5.2 (x64) Xen (xen.gz-2.6.18-8.el5)	RHEL 5.2 (Para-virtualization enabled)	4 Virtual CPUs, 4 GB RAM (pre-allocated disk in file)
RHEL 5.2 (x64) Parallels 4.0 (x64)	RHEL 5.2	4 Virtual CPUs, 4 GB RAM (Parallels Containers does not require a virtual disk)

Performance Notes

- **Parallels Containers/Xen performance “kernel per minute” can be better than native node numbers due to better compilation scalability in CTs/VMs**

- **Could not run tests for 16 Xen VMs due to memory limitations in Xen**
  - Parallels Containers could compile in parallel in both 16 & 32 containers
  - Parallels Containers compile performance decrease minimally as CT count grows
  - Xen is not stable at high VM numbers
  - We limited testing to 32 CTs/VMs due to Qualcomm density recommendations

- **Compilation time in containers shows average than 5-10% difference**
  - Compilation process not only starts at the same time in all containers, but also finishes at approximately the same time.
  - In this case kernel compilation speed (kernel per minute) can be evaluated as number_of_CTs*1/average_compile_time=3*1/7=0,42 very close to calculated 0,43

- **Compilation performance variance depends on many factors**
  - i.e. disk bottleneck (slow hardware storage), increases disk access….

Benchmark Results

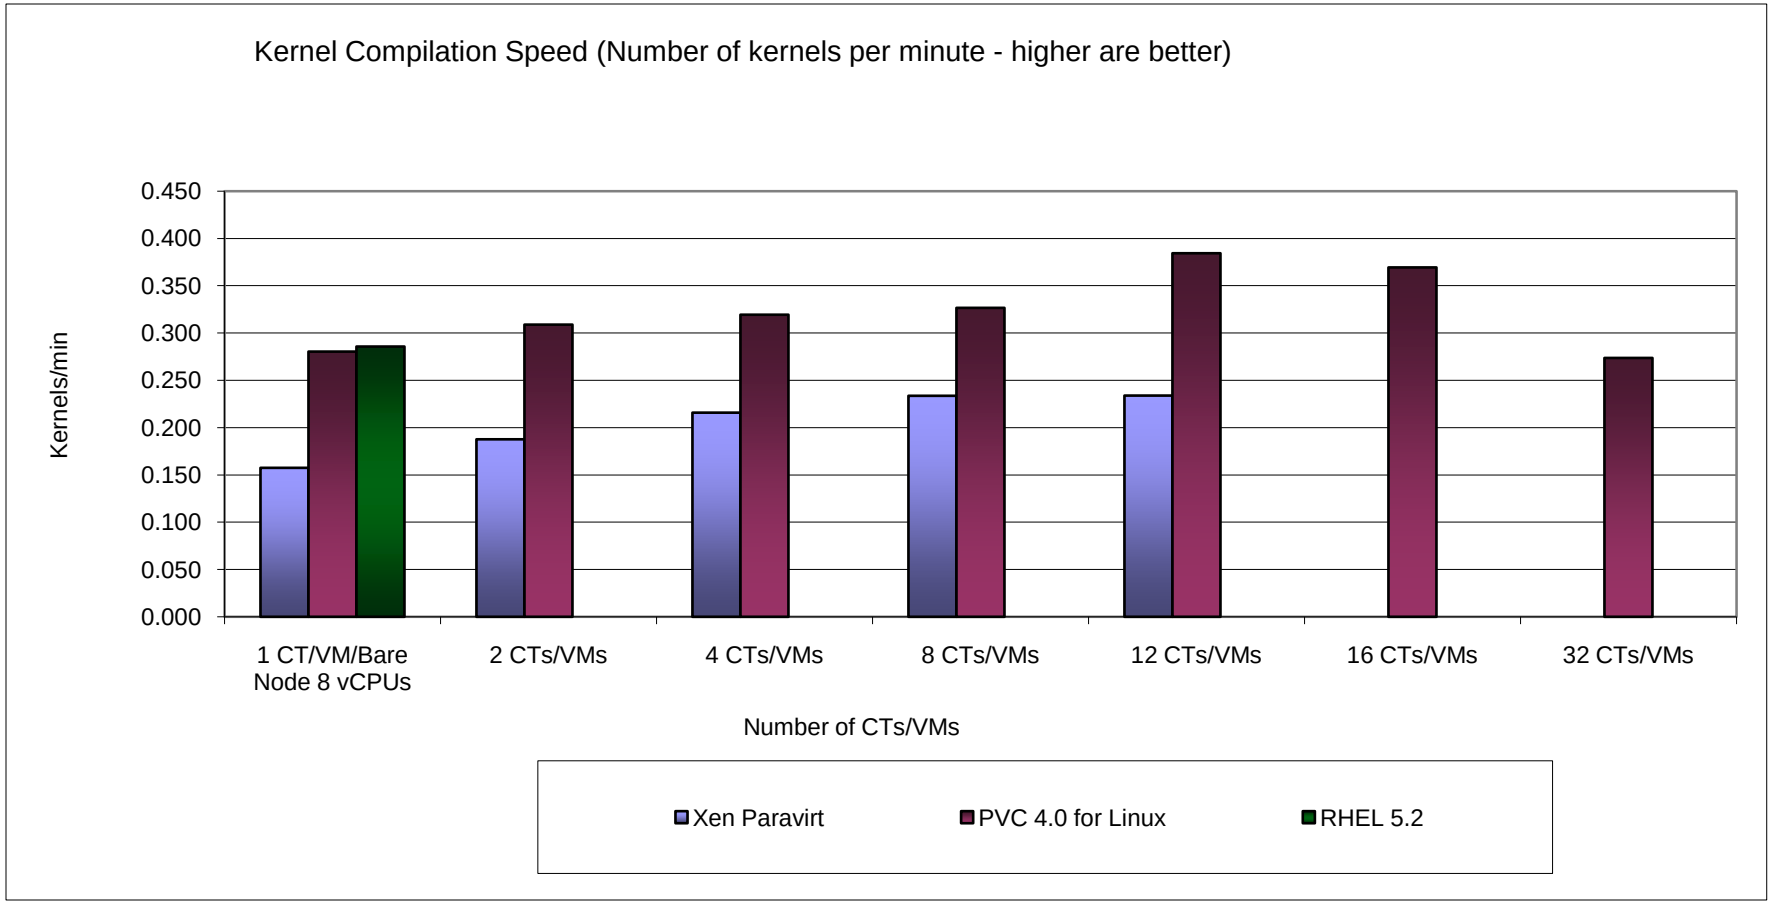

Parallels Containers achieved 40-78% better performance

Benchmark Results

Average kernel compilation time in each CT/VM

Average number of kernels per minute that can be compiled in parallel (number_of_VMs*60/(average_kernel_compile_time_in_seconds)

Percentage Parallels is faster at compiling a kernel under load vs. Xen Para-virtualization

CT / VM Count	Kernel Compilation (sec)			Compilation (kernels/min)			Parallels Performance Benefit
	Parallels	Xen Paravirt	RHEL 5.2	Parallels	Xen Paravirt	RHEL 5.2	
Bare Node			*210*			*0.286*	
1	214	381		0.280	0.157		**78%**
2	389	640		0.309	0.188		**65%**
4	752	1,112		0.319	0.216		**48%**
8	1,470	2,054		0.327	0.234		**40%**
12	1,873	3,079		0.384	0.234		**64%**
16	2,599	n/a		0.369	n/a		
32	7,016	n/a		0.274	n/a		

Parallels Containers achieved 40-78% better performance

Average number of kernels per minute

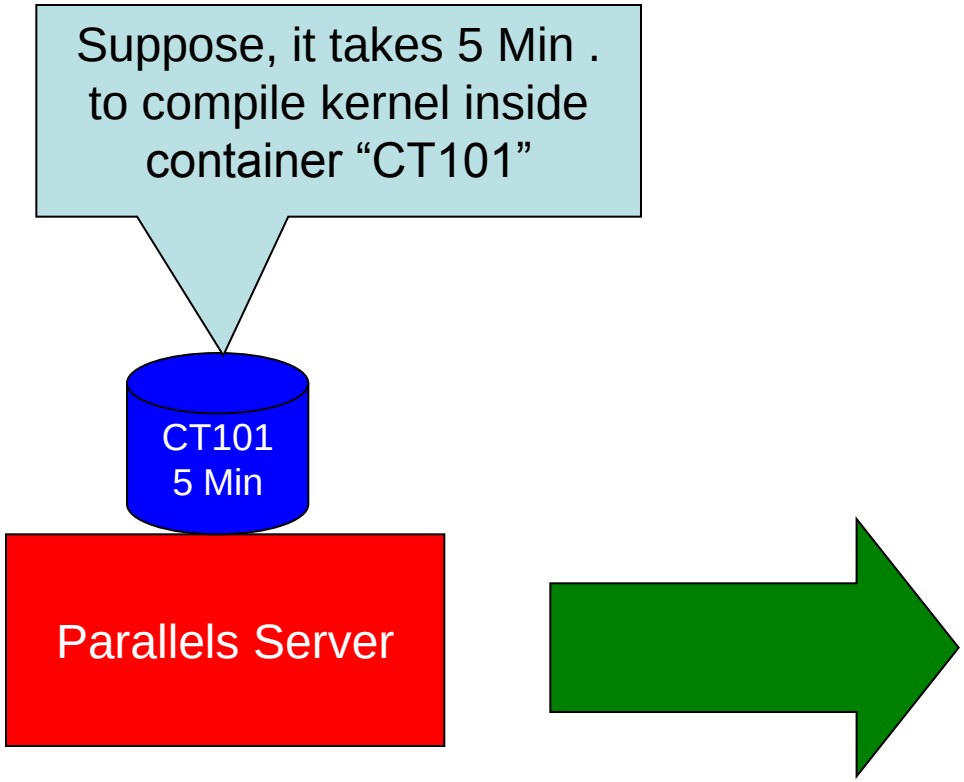

Per minute this container can compile 1/5=0.2 Kernels

Average number of kernels per minute

Suppose, it takes 7 Min . to compile kernel inside container “CT101”, 6 Min to compile kernel in “CT102” and 8 Min to compile kernel inside “CT 103”. The compilation was started at the same time in all containers.

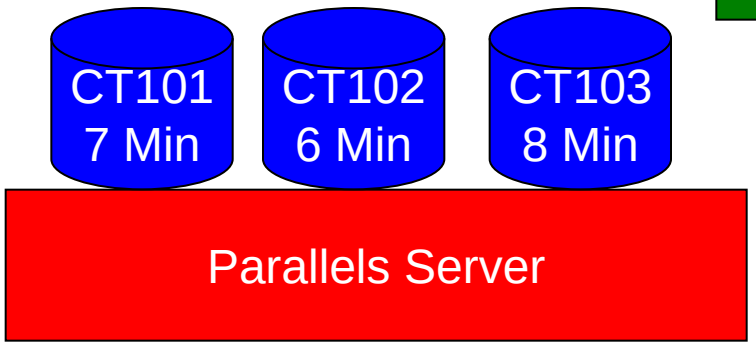

So, “CT 101” compiles 1/7 kernels per minute, “CT 102” 1/6 kernels per minute and “CT 103” 1/8 kernels per minute. And they do this compilation at the same time! **Total server** compilation speed: 1/7 + 1/6 + 1/8~**0.43** it’s >two times faster than when we use 1 CT (0,2 kernels per minute on the previous slide).

Parallels Containers Scalability on Linux – Results

- **Unique OS Virtualization approach enables neutral to positive throughput when compared bare OS baseline results across a range Linux applications**
  - Assigning one Container to each CPU Core yields the maximum results
  - For environments which need to support higher isolated workload, Parallels Containers yields solid results at two Containers per Core

- **Parallels Containers scale to enable maximum utilization on server hardware while providing complete workload isolation**
  - No limitation of resource allocation per Container → Scale up or Scale out

Summary of Parallels Elastic IT Platform

Enterprise Can Save Millions!!!

- Consolidating Servers through **Virtualization**
  - *Eliminate Underutilization*
  - *Reduce Needed Datacenter Space*
  - *Reduce Energy Costs*
  - *Reduce Storage Costs*
  - *Reduce Software Licensing Costs*

- Mass Manage through **Automation**
  - *Optimize Capacity Planning*
  - *Dynamically Allocate Resources*
  - *Quickly Mass Deploy Applications*

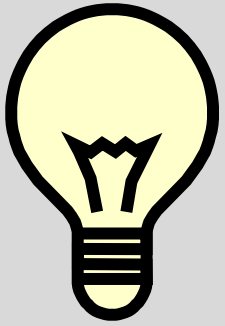

Over the next 5 years, most enterprise data centers will spend as much on energy (power and cooling) as they do on hardware infrastructure.

Gartner

Gartner, Inc. "Eight Critical Forces Shape Enterprise Data Center Strategies"

Parallels VDI: A Complete End-to-End Solution

- **The *Most* cost-effective solution to deliver server-based desktops.**
- **Thin Client/Desktop & Connection Broker**
  - Familiar and Flexible End User Experience
- **Desktops Imaged on a Server**
- **Management Toolset**

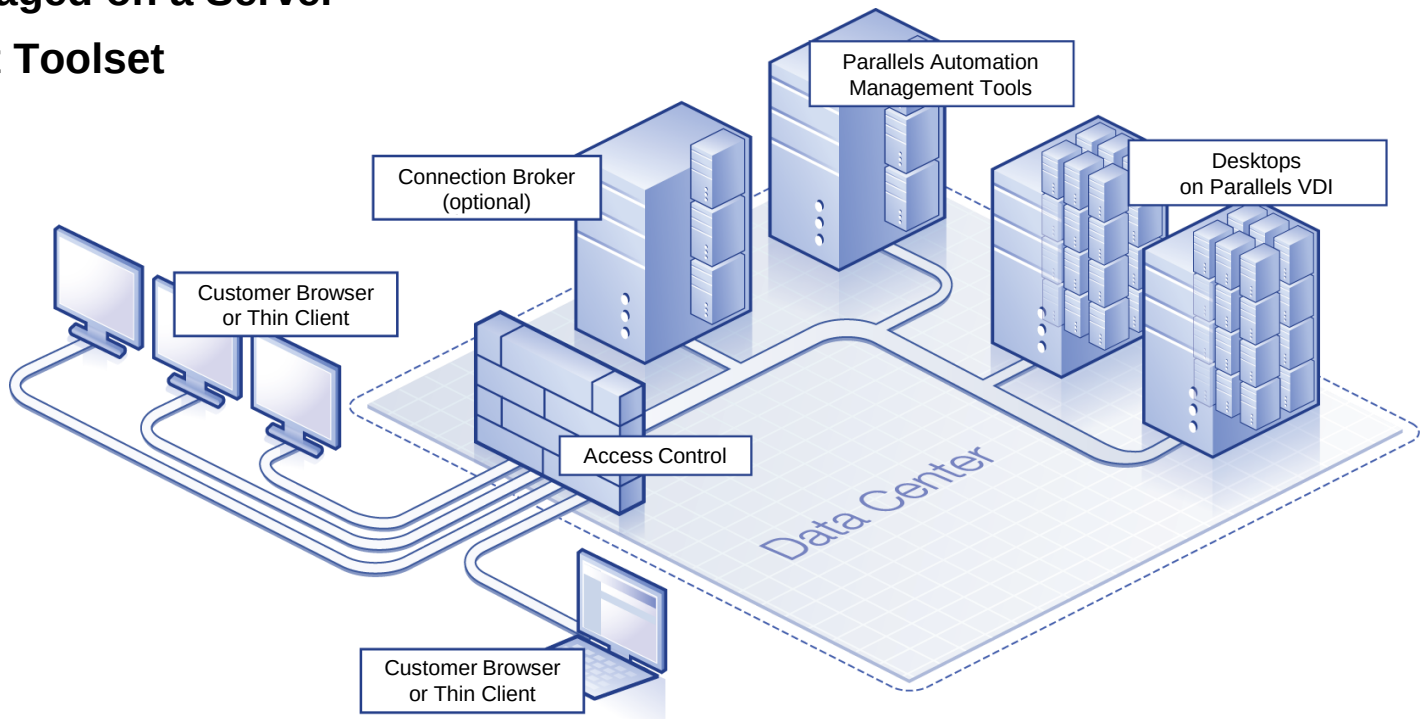

Parallels VDI Solves Desktop Challenges

Reduce Operating Costs	**Total cost of ownership reduced 70% over standalone PC environments**
Minimizing Security Risk / Meeting Compliance Guidelines	**Data centrally stored ensures security, control and reliability. Parallels VDI only solution including backup tools**
Improved Manageability Reduce Support Requirements	**Provision desktops, add/update applications, users, groups in seconds , eliminate end-user request visits. increase administrative efficiencies 70%**
Maximization of System Utilization	**Consolidation ratios 3 times higher than hypervisors results in 100's of desktops per server significantly reducing energy consumed and support costs.**

Gartner estimates about 50 million desktop users will be on VDI by 2013.

Parallels VDI: 7 X Desktops versus leading Hypervisor

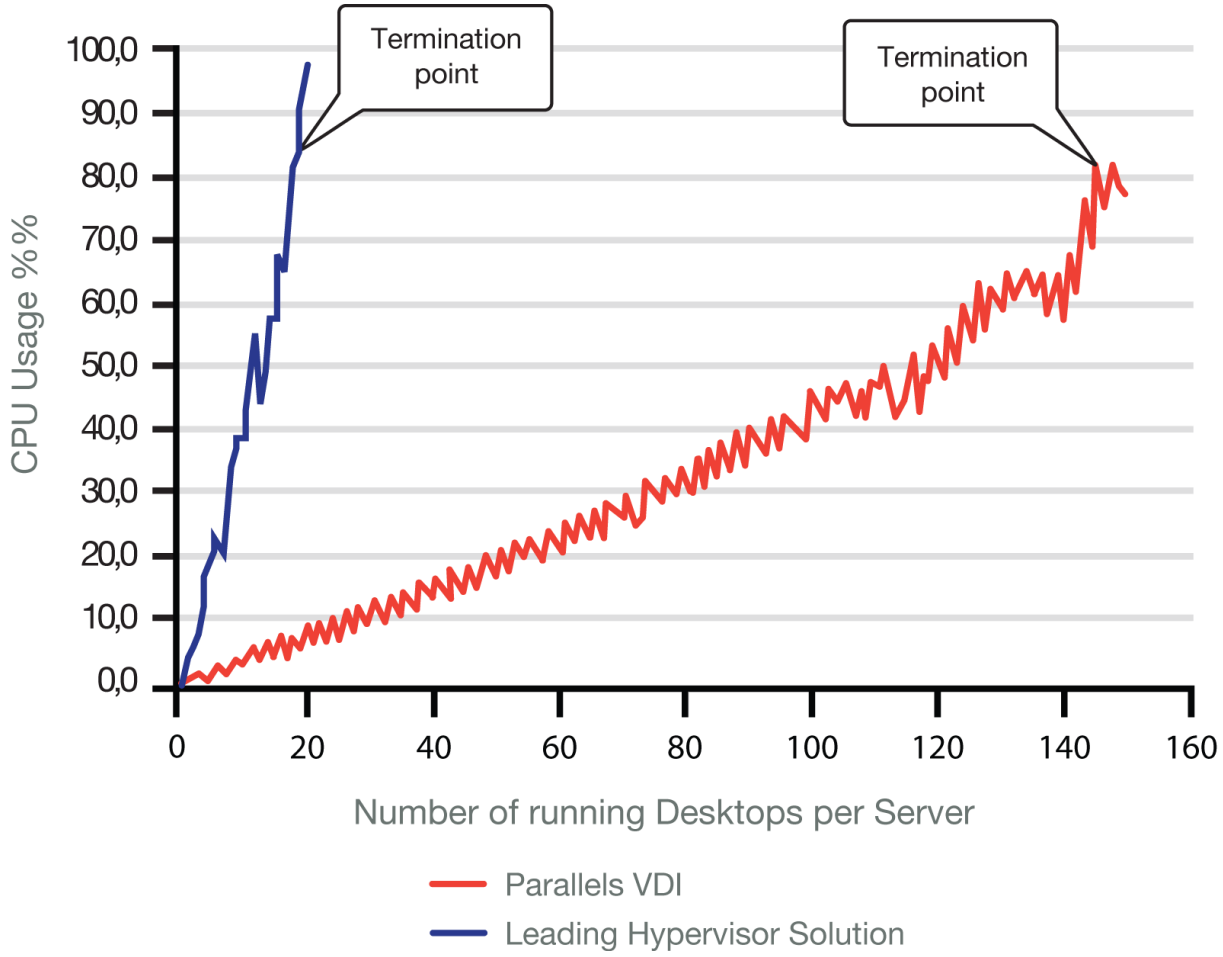

Superior Desktop-Server Consolidation Ratio: Up to 140 vs. ~20 Desktop Images

Tests run on a DL380/G5 (www.parallels.com/solution/vdi)

Parallels™

Optimized Computing™